

The Autonomy of Syntax*

David Adger

Queen Mary University of London

April 4th, 2017

1 Introduction

Chomsky's 1957 *Syntactic Structures* is, more or less throughout, an argument for the existence of syntax as both necessary for a full account of language, and as independent of other aspects of language. The book begins with arguments that the phenomenon of grammaticality cannot be explained by appeal to either probabilistic or semantic factors, and it ends with an argument that a theory of grammaticality can nonetheless provide a serious foundation for understanding questions of meaning.

This idea of the autonomy of syntax is fundamental to the generative perspective on language, but is often misunderstood. It does not entail that grammaticality is cut off from either probability or meaning. Rather it says that syntax cannot be reduced to either of these.

*Many thanks to Norbert Hornstein and Tal Linzen for comments on an earlier draft.

I think it's fair to say that the agenda laid out in *Syntactic Structures* has proved robust in both of these domains. There is no current theory of linguistics that takes human syntactic capacities to be reducible to general probabilistic relations between elements of unanalysed data, and the various research efforts into the nature of syntax have borne fruit in terms of a rich understanding of many questions of semantics that were unaskable beforehand (Portner and Partee 2002).

Beyond the claims about autonomy of syntax from statistics and semantics in the initial and final chapters, a further argument is given in the central analytical section of *Syntactic Structures*, chapters 3 through 7, for a different but equally important notion of autonomy.

This argument is based on analytical success. It says: here is a theory of some important and complicated facts about a human language (English). This language has certain properties that are best captured by a fundamentally computational model that defines an unbounded collection of hierarchically organised structures for English sentences at a number of different levels. It is therefore profitable to think of English as involving a system of that sort. Systems of this sort (finite devices that specify an infinite structured output) are generative grammars. Therefore the best theory of English incorporates a generative grammar. Further, the generative grammar given in *Syntactic Structures* for English is a more empirically successful theory than its competitors, and so the basic notion of an autonomous system of rules should be taken seriously. This is the notion of autonomy that Chomsky later explicitly sketched in 1982 (Chomsky et al. 1982): syntax as a computational system that interfaces with both semantics and phonology but whose

functioning (that is the computations that are allowed by the system) is not affected by factors external to it.

In this paper I reprise the arguments for an autonomous syntactic component given over 60 years ago in the context of more recent research, and I evaluate whether current conceptions of syntax in the Minimalist Program are consistent with the position laid out then.

2 Grammaticality and probability

Chapter 2 of *Syntactic Structures*, titled “The independence of grammar”, asks the question of how to go about separating grammatical sequences of elements from ungrammatical ones. The fundamental task the book engages with here is an explication, in theoretical terms, of an intuitive notion of grammaticality. Of course there is an existential presupposition buried in this task: that there is something in human language that the intuitive notion of grammaticality corresponds to—that is, that grammaticality is a real phenomenon. The chapter argues that there is indeed a notion of grammaticality because sentences of a language like English have properties that cannot be reduced to what is observed, and cannot be explained by what is meaningful, or by what has a high order of statistical approximation to a language.

Chomsky very briefly first makes the point that the notion of grammatical sentence cannot be identified with a particular finite collection of utterances, as any grammar has to go beyond what has been observed to new utterances that are not in a specified finite corpus. This brief comment is actually

fundamental to the argument made in chapters 3 through 7 of *Syntactic Structures*, as it makes generativity a central desideratum for a theory of grammaticality. If a grammar has to go beyond a corpus, then it is not a list, so the question arises as to what sort of a thing it is. An obvious candidate is that a grammar specifies the legitimacy of sequences of linguistic elements. But then we need some notion of ‘legitimate’. Chomsky identifies two candidates—that the sequence is some level of statistical approximation to the corpus, or that the legitimacy of the sequence is defined in terms of a notion of significance. He shows that both are inadequate, and suggests a solution that takes the grammar not to determine legitimate sequences, but rather legitimate structures. But this requires a generative grammar, defined in terms that are not reducible to statistical approximation to a language or in terms of meaningfulness.

Chomsky’s argument against defining grammaticality in terms of statistical approximations to English is fundamentally an argument against an approach to language that stems from Shannon’s work. Shannon (1948) provided a series of texts whose $(k+1)$ th symbol was determined by a probability appropriate to the preceding k symbols. If the symbols are words, then the higher order approximations begin to look more and more like English. For example, Miller and Chomsky (1963) give the following 5th order approximation (that is with quintuples of words where the frequency of the transition between one word and the next in the sequence are representative of an extant corpus of English):

- (1) road in the country was insane especially in dreary rooms where they have some books to buy for studying Greek

Chomsky argues, however, that a Shannon-style notion of statistical approximation to English for a given sequence cannot be identified with the notion of grammaticality. Using the (now) famous pair of sentences in (2), he pointed out that both would be ruled out as being equally far from a statistical approximation to English. An approach along such lines would collapse impossible sentences and sentences with extremely low probability into a single class.

- (2) a. Colorless green ideas sleep furiously.
b. Furiously sleep ideas green colorless

Later work, however, (for example Pereira 2000), has shown that one can in fact treat both sequences of words as having calculable probabilities. Pereira uses a statistical bigram model, where the probabilities are calculated by a method that includes a hidden variable encoding the class that the first word falls into, which conditions the likelihood of the second word. This allows the model to assign probabilities to pairs of words that don't exist in the training corpus. This, Pereira argues, sidesteps Chomsky's concern. Using this method, Pereira shows that the grammatical sequence in (2-a) has a much higher total probability than the ungrammatical sequence in (2-b) (setting aside, of course, cases where they are used in discussions of linguistics). Lappin and Shieber (2007) argue that Pereira's work indicates that "information theoretical approaches to modelling grammar are not vulnerable to the simple arguments for rich innate structure" widely accepted since *Syntactic Structures*.

However, noone (and I include Pereira, and Lappin and Shieber here)

thinks that a bigram model smoothed in this way could serve as a statistical model for grammaticality. Further, Chomsky’s argument here is not for “rich innate structure”. His argument is simply that the two sentences will be treated as having the same *kind* of status by a statistical model, though they behave quite differently phenomenologically. As *Syntactic Structures* points out, native speakers read the examples with different intonation and have different capacity to recall them. Miller and Isard (1963) showed that humans can more easily pick out examples like (2-a) from examples like (2-b) in noisy environments. More recent work (Pallier et al. 2011) has shown that different neural structures are recruited when sequences like (2-a) are encountered from those that are activated by ungrammatical sequences like (2-b).

Further, recent work in Artificial Language Learning provides a new argument against the idea that our knowledge of grammar is stored via knowledge of the likelihood relations between words. Culbertson and Adger (2014) presented subjects with an online task of learning an artificial language, where the word order was distinct from English, but the lexical items were identical. Subjects were presented with a noun followed by one modifier (a demonstrative, or an adjective, or a numeral). Once subjects had learned this, they were asked to translate phrases with two modifiers into this language, requiring them to extend their knowledge beyond what they had evidence for. If they were storing their knowledge of order by transitional probabilities sensitive to syntactic class, then the prediction would be that they should preserve as much of the word order of English as possible. So they should favour examples like *boxes three big* or *balls those four*, since that would preserve

the transitional probabilities between numerals and adjectives and between demonstratives and numerals. If they stored their knowledge structurally, in terms of what categories are closer to the noun, irrespective of sequence, then they should favour *boxes big three* and *balls four those*, as these preserve the closeness of the categories to the noun that we see in English. Culbertson and Adger showed that their subjects overwhelmingly preferred the latter kind of response, making the storage of ordering information as transitional probabilities between words or classes very unlikely. The online experiments have since been replicated in the lab, and are currently being extended to speakers whose first languages have other orders.

The crucial point made in *Syntactic Structures*, which remains true to this day, is that, as far as a statistical model is concerned, grammatical strings and ungrammatical strings are treated as phenomenologically the same, differing only in likelihood; as far as human interpreters and producers of language are concerned, however, they behave as though they are different kinds of thing. In order to capture this difference, a statistical model would have to somehow be mapped to a model of something which, in the end, amounts to grammaticality.

Syntactic Structures also provides a second argument against probabilistic models. This argument turns on the long distance nature of linguistic dependencies. If we examine a sentence with a relative clause attached to the subject, the question is how to capture the relationship between the grammatical number of that subject and a following auxiliary verb. Chomsky gives

(3) The man who . . . are here.

Chomsky points out that examples like (3) raise a logical problem. There is no limit to the size of the relative clause that can be inserted where . . . appears in (3), yet the sentence will remain ungrammatical nevertheless. A bigram, trigram, n-gram model, no matter how extended, is unable to capture this true fact about the grammar of English, as any value of n will be finite. Chomsky's conclusion is, again, that we need an independent model of syntactic structure to capture dependencies that are, in linear terms, unbounded.

Recent work using Recurrent Neural Networks (RNNs) which have a Long Short-Term Memory (LSTM) architecture (of the sort recently released by Google) have proved impressive in their capacity to capture long-distance statistical regularities in texts. RNNs can extract statistical regularities across vast datasets; unlike n-gram models, they do not impose any specific limit on the distance between elements in the dependency, which means that they don't suffer the same logical problem that Chomsky pointed out. With the LSTM architecture, RNNs remember information that has already been encountered. Can these address the worry laid out 60 years ago in a footnote in *Syntactic Structures*, using only statistical properties of sequences to explicate the notion of grammaticality?

Linzen et al. (2016) tested such a model on exactly the type of phenomenon Chomsky raised: number agreement. They found that a LSTM neural net could succeed on a number of cases where NPs of a different number intervened between the head of the subject NP and the agreeing verb.

- (4) a. The only championship **banners** that are currently displayed within the *building* **are** for national or NCAA Championships.
- b. Yet the **ratio** of *men* who survive to the *women* and *children* who

survive **is** not clear in this story.

Intuitively, and somewhat simplifying, these networks are tracking the likely statistical regularities between particular nouns and verbs (basically how likely it is that the noun is a subject of the verb—what linguists would think of as semantic selection) and then that information is used to key the likelihood of matching for number.

The model succeeds pretty well on most examples that Linzen et al. tested it on that involved simple subject verb agreement, but, of course, for the purposes of Chomsky’s argument, the crucial cases are where there is an unbounded dependency between the head noun and the verb. When the network was tested on relative clauses without complementizers, in examples like (5), it failed on a quarter of the cases.

- (5) The **landmarks** this *article* lists here **are** also run-of-the-mill and not notable.

Moreover, for cases where there were intervening NPs between the head noun of the subject and the agreeing verb, as in (4) and (5), the model only achieved any measure of success when it was explicitly trained on the relevant data. The training of the network effectively provided explicit information as to where the verb would be (hence what the likely syntactic boundaries were) and that the important property of the verb was its grammatical number. Because grammatical information was explicitly fed into the network, it is an unlikely model for human language acquisition. When Linzen et al. used weaker supervision methods, for example, just training the network on whether a sentence was grammatical or ungrammatical, error

rates rose substantially.

More relevant to the arguments made in *Syntactic Structures* is how the model behaved when it was unsupervised. In a pure language modelling task, with no grammatical information fed into the network, the network failed miserably, doing worse than chance. This suggests that it was actively tracking the properties of the irrelevant intervening NPs as opposed to learning anything about what the subject of the sentence is and its role in triggering agreement on the main verb. Of course, it is such a pure language modelling task that would identify the notion of grammatical sentence with probable sentence, and even with sophisticated neural net models, it seems that Chomsky’s point from six decades ago still holds strong. Human beings don’t seem to work in the way that statistical models do.

It is true that humans do, in production and in comprehension tasks, make errors somewhat like those made by Linzen et al.’s network. Bock and Miller (1991), and substantial follow-up research (Bock et al. 2001), demonstrates that people make “agreement attraction” errors, where the main verb agrees with the number of a linearly preceding noun which is not the head of the subject:

(6) The key to the cabinets are on the table

However, there are quite substantial differences between the behaviour of human subjects and the RNN models tested by Linzen et al. Whereas the RNN models show a substantial effect of whether the intervening noun is in a PP or a relative clause, humans show no such effect (Linzen et al. 2017). Wagers et al. (2009) further argue that the attraction effect can be detected in

sentence comprehension in ungrammatical, but not in grammatical sentences. This, they argue, suggests that it is a real-time processing effect, as opposed to being baked into the representation of linguistic information.

Further, the existence of agreement attraction errors actually strengthens Chomsky's argument for distinguishing grammaticality from probability. Agreement morphology is particularly fragile in language change (Ferguson 1996). If there is no independent notion of grammaticality, then the combination of the fragility of agreement in historical change, the existence in the input of agreement attraction errors like (6), and the putative probabilistic nature of the grammar, should lead us to expect the historical development of systems of agreement in natural languages that are organised via a principle of contiguity: always agree with the immediately preceding NP. This is clearly the default behaviour for the RNN network. But human languages just don't work like that. There are special cases where verbs agree with an adjacent NP, such as in first conjunct agreement (Aoun et al. 1999, Munn 1999), but no language known organizes its agreement system in the way that would make things easy for an RNN.

This is essentially the flip side of Chomsky's argument in *Syntactic Structures*, but it also makes the case for autonomy. We now have a Goldilocks argument: probabilistic models are at once too powerful and not powerful enough. On the one hand, even RNNs augmented with the LSTM architecture are incapable of learning that grammatical agreement obtains between the head of a subject and the verb, irrespective of intervening material. On the other, many statistical models can easily learn all sorts of dependencies that we simply don't find in human languages. Ergo, grammatical capacity

is a distinct system from statistical capacities.

However, one should not therefore dismiss the importance of work on probabilistic models. Chomsky's arguments are fundamentally that grammaticality cannot be identified with the probability of a sequence, so something more than that is needed. The crucial point is really about the inadequacy of sequences as a fundamental data structure for grammaticality; even when augmented with probabilistic information, they are still the wrong kind of basis for a theory language. Rather than sequences, Chomsky argues that the grammar must generate structures. Without doubt humans are sensitive to probabilities of words in contexts, and probabilistic models can provide insight into our use of language, in acquisition, in processing (Bod et al. 2003) and in broader contexts, such as sociolinguistic aspects of language (Labov 2001). However, the usefulness of probabilistic information is orthogonal to whether sequences augmented with such information are a basis for grammaticality. It transpires that they are not.

3 Grammaticality and meaning

Chapter 2 of *Syntactic Structures* also provides two brief arguments that the notion “grammatical” cannot be identified with “meaningful” or “significant”. Chomsky was reacting here to Quine's claim that linguists need to solve the “problem of defining the general notion of significant sequence.” (Quine 1953, page 53). Of course, Quine is in some sense correct—it is important for linguists to develop a theory of linguistic meaning—but Chomsky's point was that in order to address Quine's problem, we cannot assume some notion of

“significance” in advance. Rather, the way to proceed is to work out a notion of grammaticality, and from there determine a theory of meaningfulness.

The first argument Chomsky gives is again on the basis of *Colorless green ideas sleep furiously* and its reversal. While neither is “significant” in any sense that Quine would accede to, they are clearly distinct in nature. We can see this even with a more minor modification of Chomsky’s examples:

- (7) a. Colorless green ideas sleep furiously.
- b. *Colorless green ideas sleeps furiously

The difference in acceptability between the two examples here cannot be tied down to “significance” or meaningfulness in any obvious sense. It follows that we cannot identify the notion of grammatical sentence with that of significant sequence.

The other side of this argument is the contrast between cases where the meaning of a sequence of words is obvious by analogy with closely related sequences, but, irrespective of our ability to assign meaning, we still want to say that a sequence of words is ungrammatical.

- (8) a. The book seems interesting
- b. *The child seems sleeping

Since (8-b) is intuitively unacceptable as an English sentence, but perfectly interpretable on analogy with (8-a), again there is a need to appeal to grammaticality as something distinct from meaningfulness.

The conclusions of these arguments have been broadly accepted in linguistics. Newmeyer, in a series of publications (e.g. Newmeyer 1983, Newmeyer

1998), has provided a huge range of examples which bolster both lines of argument. For example, the close paraphrase relation between (9-a) and (9-b) is irrelevant to the passivizability of one and not the other:

- (9) a. Anson likes Lilly
b. Anson is fond of Lilly
- (10) a. Lilly is liked by Anson
b. *Lilly is fonded (of) by Anson

What would a theory that denies the autonomy of syntax from semantics look like and how could it tackle the kinds of problems raised by examples like these? Cognitive Grammar (Langacker 1987), and the closely related Cognitive Construction Grammar (Goldberg 2003), approach this problem by combining statistical learning methods with rich semantic representations which, though innate, are not language specific, rather they are derivative of cognition general mechanisms and representations. The approach these theories take to Chomsky's arguments are that, while grammaticality cannot be reduced to either statistical pattern matching methods or to semantics, the combination of the two is sufficient to provide an explication of grammaticality.

For example, Langacker (2008) addresses (7) by arguing that the various words in the sequence are the phonological 'poles' of symbolic units whose semantic 'poles' are interwoven via learned associations. Langacker calls these learned relationships "constructional schemas", and argues that they are abstracted from encounters of occurring expressions through a cognition general process. The reason that the (b) example in (7) is ungrammatical is

because the patterns in the sequence do not conform to any constructional schemas. In contrast, the (a) example in (7) does conform to the relevant constructional schemas. However, inconsistencies at the semantic level lead to a difficulty in assigning a coherent meaning to the whole.

We've already seen the challenges that non-adjacency in agreement raises for general pattern matching schemes: they do not learn the correct information to capture grammatical dependencies in even something as simple as subject-verb agreement. However, Langacker's system (and Goldberg's) adds semantic information to statistical processing. This semantic information is derived from the innate cognitive structure of the human mind interacting with experience, so the innate structure can be used to constrain the hypotheses that the probabilistic learning mechanism can use. Following Quine, these theories are "knowingly and cheerfully up to [their] neck[s] in innate mechanisms of learning readiness" (Quine 1969). One such innate mechanism is the undoubtedly important Figure Ground relation (Talmy 1975), which underpins the difference between subjects and objects in these theories. The Figure (which Langacker calls the Trajector) corresponds to what traditional grammar would identify as the subject, and subject-verb agreement springs off of this non-linguistic cognitive representation. This allows such theories to take the head noun in the subject to be the element which is involved in the agreement relationship—which is then just a conventionalization of the relationship between the cognitive semantic representations and the morphophonological form, with no autonomous syntactic level between them.

However, I think it's clear that Chomsky's judgment about the impossi-

bility of building a theory of syntax on semantic grounds still holds. Indeed, work in generative grammar over the decades since has revealed many reasons to understand certain semantic phenomena as supervening on syntactic structure. Without an independent syntactic representation, no amount of rich innate semantic structure is sufficient to explain how language works.

Let us take, for example, the case of the relationship between quantified NPs and pronouns that they are semantically construed with, such as classical cases of bound variable anaphora (Reinhart 1976). Adger and Svenonius (2015) use these to argue for the superiority of a structural over a semantic account.

(11) No woman denies that she has written a best selling novel.

Here the pronoun *she* can be construed as being bound by the quantifier phrase *no woman*. Certain kinds of relationship between the quantifier and the pronoun make this reading impossible.

(12) A man who no woman likes denies that she has written a best selling novel.

Here the quantified NP is embedded inside a relative clause. In such circumstances, the meaning where the quantifier binds the pronoun vanishes. This effect is cross-linguistically widespread (Déchaine and Wiltschko 2014) and experimentally robust (Kush et al. 2015). The current consensus in syntactic theory is that this phenomenon is best explained by specifying a relationship between semantic scope and binding along the following lines (Safir 2004, Barker 2012):

- (13) The Scope Generalization: For a quantifier to bind a pronoun it must scope over that pronoun.

Scope, in such theories, is constrained by syntax. In many languages, including English, a quantifier cannot scope outside of a finite clause. Since *no woman* is in a finite embedded relative clause in (13), it cannot scope at the matrix clause level and hence cannot take scope over the pronoun *she*. (13) therefore accounts for (12), when combined with a specification of what scopal freedom a language allows its quantifiers.

There is good evidence going beyond languages like English that a quantifier's ability to bind a pronoun tracks its ability to scope over the pronoun, and that scopal freedom is fixed differently in different languages, depending on the syntactic mechanisms of the language.

In Chinese, for example, the particle *dou* appears in the clause where an associated quantifier takes its scope. In (14), *dou* appears inside a relative clause, together with its associated quantifier phrase *meige ren*, 'every man'. Since *dou* marks the scope of that quantifier phrase as falling inside the relative clause, the Chinese example works just like English. The quantifier cannot scope outside of the relative clause, and a bound variable interpretation of the pronoun *ta* is impossible (Huang 1982).

- (14) meige ren dou shoudao de xin shangmian you ta taitai de
 every man DOU receive DE letter top have he wife DE
 mingzi
 name
 'The top of the letter that every man received has his wife's name
 (on it)'

However, the syntax of Chinese allows more freedom in how quantifiers scope than the syntax of English does. While in English a quantifier cannot scope outside of a relative clause, in Chinese a quantifier scopes just where *dou* tells it to. That is, *dou* marks the scope of its associated quantifier phrase, even when *dou* appears in the main sentence, rather than in the relative clause. This may be because finiteness is the crucial factor and finiteness differs in Chinese and English or it may be that the quantifier in Chinese is really *dou*, and that the apparent quantifier phrase just functions as a modifier of *dou* (extending the proposals in Shimoyama 2006). In any event, when *dou* appears in the main sentence, it marks the scope of the quantifier phrase embedded inside the relative clause as being the whole sentence, rather than just that relative clause. As predicted, the pronoun *ta* can now receive a bound variable meaning:

- (15) meige ren shoudao de xin shangmian dou you ta taitai de
 every man receive DE letter top DOU have he wife DE
 mingzi
 name
 ‘For each person, the tops of the letters he received have his wife’s
 name on them’

The Chinese facts strongly support the Scope Condition, as the pronoun in (14) and (15) can only receive a bound variable interpretation when the quantifier scopes over it, as marked by the grammatical position of *dou*. They also strongly support the notion that scope is (morpho)-syntactically determined.

The Scope Condition is empirically solid across languages where the bound variable anaphora phenomenon is detectable and it is a condition that

connects semantics to syntactic structure, and ultimately, phonological form. To the extent that it is correct, it would have to be learned, in a system like Langacker's, as a constructional schema. However, this strains credulity, as an identical abstract principle would have to be learned across a whole range of unrelated languages with quite different word order properties. The scopal behaviour of quantifiers is also not obviously detectable from surface word order properties, so it is unclear how the phonological poles of the relevant constructional schemas could ever be learned.

An alternative is to provide a solution to the phenomenon that denies the relevance of the Scope Condition and builds an explanation of the facts on purely semantic grounds. This tack is taken by van Hoek (1996), who argues that whether a pronoun can be bound is dependent on the salience or prominence of the quantifier. Van Hoek defines salience in terms of a Figure Ground relationship, where the Figure is always more salient and can act as the "binder". Adger and Svenonius (2015) argue, however, that this approach does not fare well empirically. Talmy (2000) defines Figure Ground in the following way:

- (16) a. The Figure is a moving or conceptually movable entity whose path, site, or orientation is conceived as a variable, the particular value of which is the relevant issue
- b. The Ground is a reference entity, one that has a stationary setting relative to a reference frame, with respect to which the Figure's path, site, or orientation is characterized

Adger and Svenonius give numerous cases where the subject of a sentence

is the Ground, rather than the Figure but, counter to the expectations of a theory like van Hoek's, this does not impact on the distribution of bound variable anaphora. For example, the verb *contain*, by definition, has a Figure as object and Ground as subject. If the Figure is always salient, van Hoek's system incorrectly predicts the wrong binding possibilities:

- (17) a. Each book contains its author's biography as an initial chapter.
b. *Its initial chapter contains a synopsis of each book.

Examples of this kind, with the verb *contain*, also raise problems for approaches like that of MacWhinney (1977) et seq., where the notion of 'perspective' is used as a constraint on how various kinds of dependencies are learned. MacWhinney uses perspective to provide a non-syntactic, cognitive grounding to effects that are handled in generative grammar through structural command. For MacWhinney, linear and semantic factors intertwine with processing effects to determine which NP is prominent in the sentence for various apparently grammatical effects. Typically, the perspective holder is an initial NP. However, to handle cases like (17), some special statement about English needs to be made (essentially that the structural subject always ends up being the holder of the perspective irrespective of semantic considerations), which obviously recreates a space for autonomous structure. Languages like Palauan are also problematic for such an approach, which favours a semantic and linear explanation over a structural one. Palauan has VOS order, but disallows binding from the initial object, even when it is human, to the following subject (Georgopoulos 1989, p202):¹

¹The particle *a* marks NPs and is glossed, P glosses a preposition-like element that,

- (18) *temengull er a rebek el 'ad a retonari er tir
 3pl.respect P every person neighbours P their
for 'Their_i neighbours respect everyone_i'

In (18), the universal quantifier *rebek el 'ad* is the object of the verb, but, in MacWhinney's approach it is an Initial NP (as it is processed first) and it is human, and so should serve as a perspectival centre. However, (18) shows that it still cannot bind the pronoun in the subject. This is straightforward to capture in a structural account, which constrains the scope of the quantifier to its c-command domain. Since the object is inside the VP, it cannot c-command the subject. Such an analysis makes the further prediction that if we can move the object to a higher position, it will be able to c-command the subject. This prediction is borne out, but is difficult to handle in a perspective based account, as no perspective changes².

- (19) rebek el 'ad a longull er tir a retonari er
 every person 3.respect P them neighbours P their
 tir
for 'Their_i neighbours respect everyone_i'

When the object is placed to the left of the verb as a topic, leaving behind a resumptive pronoun, it can now bind into the subject, so that the reading is 'For every person, their neighbours respect them.' In MacWhinney's system, topics cannot act as binders, as they are semantic scene setters, as opposed

roughly, appears with human/specific object and possessive NPs.

²For completeness sake, I should note that a subject cannot bind into the object in Palauan either, unless it is moved rightwards. Georgopoulos takes this to show that both structural and linear factors are relevant to the explanation. Of course, this still implies that one needs a structural account.

to perspectival centres.

I have chosen here to look at two of the best developed accounts I know of which attempt to tackle binding effects in frameworks that do not have an autonomous syntactic component. These are serious attempts to address phenomena that have been discovered in generative grammar, and neither, I think, is successful. The need for an independent notion of grammaticality, and hence a theory of that phenomenon, stands firm.

4 (Some) semantics supervenes on syntax

Chomsky returns to the question of the independence of grammar in Chapters 8 and 9, which argue for an architecture where certain types of meaning supervenes on grammar as opposed to underpins it. Chapter 8 on “The explanatory power of linguistic theory” is about how a grammar which makes no reference to meaning in the rules it uses to define the grammaticality of sentences can nonetheless provide explanations of certain apparently semantic facts: why certain sentences are ambiguous while other similar ones are not, why certain collections of sentences whose members are quite distinct nevertheless share a core of meaning. Theories which capture such facts should, Chomsky suggests, be highly valued:

...we should like the syntactic framework of the language that is isolated and exhibited by the grammar to be able to support semantic description, and we shall naturally rate more highly a theory of formal structure that leads to grammars that meet this requirement more fully *Syntactic Structures*, pp102.

Chapter 9 "Syntax and Semantics" uses these results to argue that the fundamental question about the relationship between syntax and semantics is not whether semantic properties enter into syntax, but rather how the independently motivated syntactic devices are put to use for semantic purposes. As Chomsky puts it in the concluding summary chapter

More generally, it appears that the notion of "understanding a sentence" must be partially analyzed in grammatical terms. To understand a sentence it is necessary (though not, of course, sufficient) to reconstruct its representation on each level, including the transformational level where the kernel sentences underlying a given sentence can be thought of, in a sense, as the 'elementary content elements' out of which this sentence is constructed. In other words, one result of the formal study of grammatical structure is that a syntactic framework is brought to light which can support semantic analysis. Description of meaning can profitably refer to this underlying syntactic framework, although systematic semantic considerations are apparently not helpful in determining it in the first place. *Syntactic Structures*, pp102.

We have come a long way in providing a grammatical basis for meaning. Much of the last half century in formal linguistics has been focussed on how syntactic form feeds meaning, and a great deal is now known about this. In addition to providing a basis for understanding ambiguity, synonymy and entailment, which were the notions under the spotlight in *Syntactic Structures*, work in formal semantics following on from Montague's work (Thomason 1974), with its focus on truth conditions, brought to light new methods for

connecting syntax and certain parts of the theory of meaning. Emmon Bach, in his *Informal Lectures on Formal Semantics*, frames the relationship between generative syntax and formal semantics elegantly: while Chomsky put forward the hypothesis that natural languages could be described as formal systems, Montague proposed that they could be described as compositionally interpreted formal systems (Bach 1989). The interaction between these two theses led to an explosion of research that showed how seemingly intractable puzzles in the meaning of languages could be addressed. For example, *Syntactic Structures* had proposed that the kernel sentences could be thought of as ‘elementary content elements’, but this conception raised some serious puzzles about how to calculate the meanings of combinations of these elements. Many of these puzzles could be given elegant solutions by using techniques from formal semantics that amounted to treating these elementary elements as having variables in them that could then be connected with other structures in more complex sentences (Hall Partee 1975). The “striking correspondences between the structures and elements that are discovered in formal, grammatical analysis and specific semantic functions” (*Syntactic Structures*, p101), have, over the decades, become more striking and more compelling.

The success in building a theory of meaning that supervenes on a generative syntactic system might seem, at first blush, to raise the question of the autonomy of this system. In *Syntactic Structures* various grammatical elements were argued to lack meaning, including the infinitival marker *to* and the dummy support verb *do*. However, Stowell (1982) showed that *to* does carry some meaning (a kind of modal interpretation), while most current

theories of *do*-support take it to involve an extra-syntactic morphological rule (Embick and Noyer 1999). Indeed, if, as Berwick and Chomsky (2016) propose, the computational system of human language (the autonomous syntactic system) is a perfect solution to interfacing with the semantic systems of thought, one might expect a kind of isomorphism in this area.

Something along these lines is, if not a consensus, at least a major research area in the field today. There is wide, though not universal, acceptance of Borer’s arguments that syntactic information should be taken out of lexical items and that a great deal of meaning should be stated in terms of the structure, as opposed to the content of lexical items (the exoskeletal mode of explanation—Borer 2003). Papers in theoretical syntax proposing various functional categories with syntactic effects typically (though not always) also provide a way of compositionally calculating their semantic effects. Chomsky’s notion of Full Interpretation, its implementation via interpretable vs uninterpretable features, and the idea that syntactic dependencies have the consequence of eliminating the latter, lead to syntactic representations which are isomorphic with certain aspects of meaning (Chomsky 1995). In what sense, then, is the system autonomous?

One area of autonomy in such a system might be that the organization of functional categories with respect to each other in an extended projection (Grimshaw 2005) must be given as an independent syntactic stipulation, disconnected from semantics (Cinque 2013a). I think that this is a currently a minority perspective. Though little is known about how the categories in an extended projection are motivated by conceptual factors, there have been significant proposals. For example, following, ultimately, the lead of

Ernst (1998), Ramchand and Svenonius (2014) propose that the reason that C contains T, which in turn contains v, is due to a kind of semantically based mereology of propositions, situations and events. Particular categories in the clause function to transition between these aspects of the semantics of a proposition. In such a proposal the building up of syntactic structure is simultaneously the building up of some aspects of semantic structure.

There are also other ways of conceiving of the reasons for why elements in an extended projection appear where they do, however, that sustain the notion of an autonomous structure building system interacting intimately with systems of thought and meaning, while being quite distinct from them. Take an example like the following:

- (20) a. Those three green balls
b. *Those green three balls

As is well known, the order of the demonstrative, numeral and descriptive adjective in a noun phrase follow quite specific typological patterns arguing for a hierarchy where the adjective occurs closest to the noun, the numeral occurs further away and the demonstrative is most distant (Greenberg 1963, Cinque 2005). Why should this be? It seems implausible for this phenomenon to appeal to a mereological semantic structure. I'd like to propose a different way of thinking about this that relies on the way that a purely autonomous syntax interfaces with the systems of thought.

Imagine we have a bowl which has red and green ping pong balls in it. Assume a task (a non-linguistic task) which is to identify a particular group of three green balls. Two computations will allow success in this task:

- (21)
 - a. select all the green balls
 - b. take all subsets of three of the output of (a)
 - c. identify one such subset.

- (22)
 - a. take all subsets of three balls
 - b. for each subset, select only those that have green balls in them
 - c. identify one such subset

Both of these computations achieve the desired result. However, there is clearly a difference in the complexity of each. The second computation requires holding in memory a multidimensional array of all the subsets of three balls, and then computing which of these subsets involve only green balls. The second simply separates out all the green balls, and then takes a much smaller partitioning of these into subsets involving three. So applying the semantic function of colour before that of counting is a less resource intensive computation. Of course, this kind of computation is not specific to colour—the same argument can be made for many of the kinds of properties of items that are encoded by intersective and subsective adjectives.

If such an approach can be generalized, then there is no need to fix the order of adjectival vs. numeral modifiers in the noun phrase as part of an autonomous system. It is the interface between a computational system that delivers a hierarchy, and the use to which that system is put in an independent computational task of identifying referents, plus a principle that favours systems that minimize computation, that leads to the final organization. The syntax reifies the simpler computation via a hierarchy of categories. This means that one need not stipulate the order in UG, nor, in fact, de-

rive the order from the input. The content and hierarchical sequence of the elements in the syntax is delivered by the interface between two distinct systems. This can take place over developmental timescales, and is, of course, likely to be reinforced by the linguistic input, though not determined by it. Orders that are not isomorphic to the easiest computations are allowed by UG, but are pruned away during development because the system ossifies the simpler computation. Such an explanation relies on a generative system that provides the structure which the semantic systems fill with content.

The full ordering of the content of elements in a syntactic hierarchy presumably involves a multiplicity of sub ordering effects, some due to differences in what variable is being elaborated as in Ramchand and Svenonius's proposal, others, if my sketch of an approach to the noun phrase is correct, due to an overall minimizing of the computation of the use of the structure in referring, describing, presenting etc. In this approach, the job of the core syntactic principles is to create structures which have an unbounded hierarchical depth and which are composed of discrete elements combined in particular ways. But the job of populating these structures with content is delegated to how they interface with other systems.

How does this impact on the autonomy thesis? If all of the content of the categories that syntax works with are co-opted from other systems of thought, to what extent is syntax autonomous.

The answer to this involves recognizing two distinct areas where autonomy is relevant. One is related to the point made in the body of *Syntactic Structures*, that a generative system is required for a good explanation of syntactic phenomena, and that that system operates in a particular way,

manipulating categories via their form, and not their meaning. This is autonomy of structure building.

In *Syntactic Structures* the structure building system is intimately tied to the system of categories, but that is clearly not necessary. Chomsky (1970) largely separates off the categories from the component of the grammar that licenses structure (X-bar theory), and recent approaches to bare phrase structure do this even more radically, by taking Merge to build uncategorized structures, with a distinct labelling algorithm doing the job of providing the categories (Chomsky 2013). Adger (2013) goes further, labelling structures exocentrically via independently specified extended projection information (assumed to be derivable in the ways just discussed). This theoretical bifurcation opens up the second area of autonomy: autonomy of content. To what extent are the categories involved in labelling autonomously generated syntactic structures themselves specialized to syntax.

It is certainly the case that not every conceptual category is used in grammar. Cross-linguistic work in typological linguistics as well as in generative grammar identifies, over and over again, the same set of grammatically active categories. This is recognized by functional theorists as well as generative ones. For example Talmy (1985) remarks

There are many characteristics of an event's participants that are not marked anywhere in the verb complex, even though they seem as reasonable (from an a priori perspective) as the qualities that are marked. Thus, while an argument's numerosity and distribution can be marked, there will be no marking for its color or whether it has a symmetrical arrangement, even though these

very qualities are important in other cognitive systems, such as visual perception? (p.134).

Similarly, many languages have miratives, where a verbal form grammatically expresses that the speaker's attitude to the proposition they are making is one of surprise, but no language has a 'solicitative', marking that the speaker is worried about the situation. Cinque (2013b) provides a range of similar examples and asks the crucial question: "To say that the external and internal temporal constituency of an event (tense and aspect) or the attitude of the speaker toward the truth of the proposition (mood) are cognitively salient is beside the point. The question remains why these and only these cognitive distinctions are encoded grammatically in natural languages out of the many other salient ones." His conclusion is that UG determines what conceptual categories are grammatically encoded.

On similar lines, Hale (1986) argues that it is possible to have complete control of the grammar and phonology of a language without knowing the smallest thing about the worldview of the speakers of the language. Cultural concepts that may strongly influence how speakers of a language live their lives every day are not coopted into grammar. The way that a culture dissects the world can be embedded in the words of a language, in the myths told in the language, in idioms and turns of phrase, but grammar, and phonology, are disconnected from worldview. Hale discusses the philosophy of the Warlpiri aboriginal people of the central western parts of Australia's Northern Territory and identifies a theme, which he terms 'the logic of eternity'. This is a worldview that sees time in cyclical, rather than linear, terms and appears in many ways in how the Warlpiri conceive of objects in the

world. For example, entities are thought of as persisting, even when they are transformed: an entity is just the current incarnation of something that is eternal. When someone makes something, they are not creating it new, but rather simply altering the stuff of the world to create a manifestation of an eternal object. Hale argues that this world view can be seen in the ritual and totemic practices of the Warlpiri, in how they organise their kinship system, and in the vocabulary for words that loosely correspond to English “make”. The meanings of these words treat the act of making as the act of transforming: *ngurrju-ma-ni*, means to ‘perfect or fix’, *yirra-rni*, meaning to ‘put something in place’, and various verbs of impact, such as *paka-rni*, to ‘chop’, *panti-rni* to ‘pierce or gouge’ and *jarnti-rni* to ‘trim or sculpt’. Manufacture, or making, in Warlpiri vocabulary is coded not as creation, but as change. Hale points out that these aspects of Warlpiri culture that express the world view of the logic of eternity never seem to be at the level of grammar. He contrasts this with the kinds of concepts that often do appear in the grammar of the world’s languages: cause, tense, evidence, and so on. This makes the same point as Talmy and Cinque: of all the possible concepts that humans have, only a restricted set of these appear in grammar. Something restricts which concepts human grammar co-opts from human thought. Whatever that is, it’s another instance of grammatical autonomy.

This second autonomy thesis opens up a new research programme: what are the co-optable concepts? Why are these the co-optable ones? What principles organise these into syntactic hierarchy? How are they ‘digitized’ by the syntactic system so as to enter into syntactic computations?

Rather than threatening the autonomy of syntax, grammar focussed in-

vestigations of questions of meaning have strengthened and clarified the issue of the autonomy of syntax, as well as opening up new research questions, showing just how fertile the core set of ideas, laid down in *Syntactic Structures* 60 years ago, remain.

References

- Adger, David. 2013. *A Syntax of Substance*. Cambridge, MA: MIT Press.
- Adger, David and Svenonius, Peter. 2015. Linguistic explanation and domain specialization: a case study in bound variable anaphora. *Frontiers in psychology* 6.
- Aoun, Joseph, Benmamoun, Elabbas, and Sportiche, Dominique. 1999. Further remarks on first conjunct agreement. *Linguistic Inquiry* 30:669–681.
- Bach, Emmon W. 1989. *Informal lectures on formal semantics*. Albany, N.Y.: State University of New York Press.
- Barker, Chris. 2012. Quantificational binding does not require c-command. *Linguistic Inquiry* 43:614–633.
- Berwick, Robert C. and Chomsky, Noam. 2016. *Why Only Us?*. Cambridge, MA: MIT Press.
- Bock, Kathryn, Eberhard, Kathleen M., Cutting, J.Cooper, Meyer, Antje S., and Schriefers, Herbert. 2001. Some attractions of verb agreement. *Cognitive Psychology* 43:83 – 128.

- Bock, Kathryn and Miller, Carol A. 1991. Broken agreement. *Cognitive psychology* 23:45–93.
- Bod, Rens, Hay, Jennifer, and Jannedy, Stefanie. 2003. *Probabilistic Linguistics*. Cambridge, Massachusetts: MIT Press.
- Borer, Hagit. 2003. Exo-skeletal vs. endo-skeletal explanations: Syntactic projections and the lexicon. In John Moore and Maria Polinsky, eds., *The Nature of Explanation in Linguistic Theory*, 31–67, Stanford, CA.: CSLI Publications.
- Chomsky, Noam. 1957. *Syntactic Structures*. The Hague: Mouton.
- Chomsky, Noam. 1970. Remarks on nominalization. In R. A. Jacobs and P. S. Rosenbaum, eds., *Readings in English Transformational Grammar*, 184–221, Waltham, MA: Ginn-Blaisdell.
- Chomsky, Noam. 1995. *The Minimalist Program*. Cambridge, MA: MIT Press.
- Chomsky, Noam. 2013. Problems of projection. *Lingua* 130:33 – 49.
- Chomsky, Noam, Huybregts, Riny, and van Riemsdijk, Henk C. 1982. *Noam Chomsky on the generative enterprise: A discussion with Riny Huybregts and Henk van Riemsdijk*. Foris Pubns USA.
- Cinque, Guglielmo. 2005. Deriving Greenberg’s Universal 20 and its exceptions. *Linguistic Inquiry* 36:315–332.
- Cinque, Guglielmo. 2013a. Cognition, universal grammar, and typological generalizations. *Lingua* 130:50–65.

- Cinque, Guglielmo. 2013b. Word order typology: A change in perspective. In Theresa Biberauer and Michelle Sheehan, eds., *Theoretical Approaches to Disharmonic Word Order*, 47–73, Oxford University Press.
- Culbertson, Jennifer and Adger, David. 2014. Language learners privilege structured meaning over surface frequency. *Proceedings of the National Academy of Sciences* 111:5842–5847.
- Déchaine, Rose-Marie and Wiltschko, Martina. 2014. Bound variable anaphora, lingbuzz.
- Embick, David and Noyer, Rolf. 1999. Locality in post-syntactic operations. In *Papers in Morphology and Syntax: cycle two*, volume 34, 265–317, MIT Working Papers in Linguistics.
- Ernst, Thomas. 1998. The scopal basis of adverb licensing. In Pius N. Tamanji and Kiyomi Kusumoto, eds., *Proceedings of the North East Linguistic Society*, 127–142, University of Toronto: Graduate Linguistic Student Association.
- Ferguson, Charles. 1996. Variation and drift: Loss of agreement in Germanic. In Gregory Guy, Crawford Feagin, Deborah Schiffrin, and John Baugh, eds., *Towards a Social Science of Language: Papers in honor of William Labov. Volume 1, Variation and Change in Language and Society*, 173–198, Amsterdam: John Benjamins.
- Georgopoulos, Carol. 1989. Syntactic variables: Resumptive pronouns and A'-binding in Palauan. Ph.D. thesis, University of Utah.

- Goldberg, Adele E. 2003. Constructions: a new theoretical approach to language. *Trends in Cognitive Sciences* 7:219–224.
- Greenberg, Joseph H. 1963. Some universals of grammar with particular reference to the order of meaningful elements. In Joseph Greenberg, ed., *Universals of Language*, 73–113, Cambridge: MIT Press.
- Grimshaw, Jane. 2005. *Words and Structure*. Stanford, California: CSLI Publications.
- Hale, Kenneth. 1986. Notes on world view and semantic categories: Some warlpiri examples. *Features and projections* 25:233–254.
- Hall Partee, Barbara. 1975. Deletion and variable binding. In E. L. Keenan, ed., *Formal Semantics of Natural Language*, 16–34, Cambridge: Cambridge University Press.
- Huang, James. 1982. Logical relations in Chinese and the theory of grammar. Ph.D. thesis, MIT, Cambridge, Massachusetts.
- Kush, Dave, Lidz, Jeffrey, and Phillips, Colin. 2015. Relation-sensitive retrieval: Evidence from bound variable pronouns. *Journal of Memory and Language* 82:18–40.
- Labov, William. 2001. *Principles of Linguistics Change. Volume 2: Social Factors*. Oxford: Blackwell.
- Langacker, Ronald W. 1987. *Foundations of cognitive grammar: Theoretical prerequisites*, volume 1. Stanford university press.

- Langacker, Ronald W. 2008. *Cognitive grammar: A basic introduction*. Oxford University Press.
- Lappin, Shalom and Shieber, Stuart M. 2007. Machine learning theory and practice as a source of insight into universal grammar. *Journal of Linguistics* 43:393–427.
- Linzen, Tal, Dupoux, Emmanuel, and Goldberg, Yoav. 2016. Assessing the ability of LSTMs to learn syntax-sensitive dependencies. *Transactions of the Association for Computational Linguistics* 4:521–535.
- Linzen, Tal, Goldberg, Yoav, and Dupoux, Emmanuel. 2017. Agreement attraction errors in neural networks. In *Poster presented at CUNY Conference on Human Sentence Processing*.
- MacWhinney, Brian. 1977. Starting points. *Language* 152–168.
- Miller, George A. and Chomsky, Noam. 1963. Finitary models of language users. In D. Luce, ed., *Handbook of Mathematical Psychology*, 2–419, John Wiley & Sons.
- Miller, George A and Isard, Stephen. 1963. Some perceptual consequences of linguistic rules. *Journal of Verbal Learning and Verbal Behavior* 2:217–228.
- Munn, Alan. 1999. First conjunct agreement: Against a clausal analysis. *Linguistic Inquiry* 30:643–668.
- Newmeyer, Frederick. 1998. *Language Form and Language Function*. Cambridge, MA: MIT Press.

- Newmeyer, Frederick J. 1983. *Grammatical theory: Its limits and its possibilities*. University of Chicago Press.
- Pallier, Christophe, Devauchelle, Anne-Dominique, and Dehaene, Stanislas. 2011. Cortical representation of the constituent structure of sentences. *Proceedings of the National Academy of Sciences* 108:2522–2527.
- Pereira, Fernando. 2000. Formal grammar and information theory: together again? *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 358:1239–1253.
- Portner, Paul and Partee, Barbara. 2002. *Formal Semantics: The Essential Readings*. Oxford: Wiley-Blackwell.
- Quine, Willard Van Orman. 1953. *From a logical point of view: 9 logico-philosophical essays*. Cambridge MA: Harvard University Press.
- Quine, Willard Van Orman. 1969. Linguistics and philosophy. In Sidney Hook, ed., *Language and Philosophy*, 95–96, New York: University Press.
- Ramchand, Gillian and Svenonius, Peter. 2014. Deriving the functional hierarchy. *Language Sciences* 46:152–174.
- Reinhart, T. 1976. The syntactic domain of anaphora. Ph.D. thesis, MIT, Cambridge, MA.
- Safir, Ken. 2004. *The syntax of anaphora*.. Oxford, UK: Oxford University Press.
- Shannon, C. E. 1948. A mathematical theory of communication. *Bell system technical journal* 27.

- Shimoyama, Junko. 2006. Indeterminate phrase quantification in Japanese. *Natural Language Semantics* 139–173:2.
- Stowell, Tim. 1982. The tense of infinitives. *Linguistic Inquiry* 13:561–570.
- Talmy, Leonard. 1975. Figure and ground in complex sentences. In *Proceedings of the First Annual Meeting of the Berkeley Linguistics Society*, volume 1, 419–430.
- Talmy, Leonard. 1985. Lexicalization patterns: semantic structure in lexical forms. In T. Shopen, ed., *Language Typology and Syntactic Description III: Grammatical Categories and the Lexicon*, Cambridge, UK: Cambridge University Press.
- Talmy, Leonard. 2000. *Towards a Cognitive Semantics*. Cambridge, MA: MIT Press.
- Thomason, Richmond H., ed. 1974. *Formal Philosophy: Selected Papers of Richard Montague*. New York: Yale University Press.
- van Hoek, Karen. 1996. A cognitive grammar account of bound anaphora. In Eugene Casad, ed., *Cognitive Linguistics in the Redwoods: The Expansion of a New Paradigm in Linguistics*, 753–792, Berlin: De Gruyter.
- Wagers, Matthew W, Lau, Ellen F, and Phillips, Colin. 2009. Agreement attraction in comprehension: Representations and processes. *Journal of Memory and Language* 61:206–237.